

REMARKS

This amendment is being filed in response to the Office Action having a mailing date of July 12, 2007. Various claims are amended as shown. New claims 28-31 are added. No new matter has been added. With this amendment, claims 1-31 are pending in the application.

I. Rejections under 35 U.S.C. § 101

The present Office Action rejected claims 13-16 under 35 U.S.C. § 101. Claims 13-16 are amended as shown to address this rejection, and therefore, it is kindly requested that said rejection of claims 13-16 be withdrawn.

II. Discussion of the claims and cited references

Claims 1, 3-6, 9-15, 17-20, 22, 24, and 27 were rejected under 35 U.S.C. § 102(e) as being anticipated by Crow (U.S. Patent No. 6,944,672). Claims 2 and 26 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Crow in view of Marleux (U.S. Patent No. 7,089,486). Claims 7, 16, and 21 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Crow in view of Viswanath (U.S. Patent No. 6,798,788). Claims 8, 23, and 25 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Crow in view of Basso (U.S. Patent No. 7,065,086). For the reasons set forth below, these rejections are respectfully traversed.

1. Discussion of independent claim 1

Independent claim 1 as amended recites, *inter alia*, that the “processing” of the head fragment is “to determine said routing information, including a destination address for said head fragment” and that the updating of the non-head fragments include “applying said determined routing information that includes said destination address to said received non-head fragment and to said any stored corresponding non-head fragment.” It is respectfully submitted that these limitations are not met by Crow, whether singly or in combination with the other cited references.

Specifically, Crow repeats his redundant processing/translation on all of his fragments in order to repeatedly determine routing information for each and every fragment,

rather applying results of previously performed processing that has already determined the routing information (including the destination address) for the head fragment, thereby avoiding the need for redundant translation/processing of the non-head fragments. Crow discloses a technique where he performs address translation on a “primary fragment” by using the header information in the primary packet to identify a translation entry 90 that specifies a destination address where the primary fragment should be routed. A translation engine 80 then uses the identified translation entry 90 to translate an address of the primary packet from one address to another (destination) address. Crow describes this process in his column 6, lines 26-34 reproduced below:

“At step 110, the translation engine 80 identifies a translation entry 90 in the translation table 82 for the primary fragment 32 using the IP and transport header information 40 and 42. At step 112, the translation engine 80 translates addressing information for the primary fragment 32 using the identified translation entry 90. At this point, translation for the primary fragment 32 is complete and the primary fragment 32 may be directed to the appropriate receiver.”

Next, Crow creates a “fragment-context” 92 for the identified translation entry 90. The fragment-context 92 operates as a piece of data to associate a “secondary fragment” to the translation entry 90 that was previously used for the primary fragment. Crow describes this process in his column 6, lines 35-52 reproduced below (emphasis ours):

“Proceeding to step 114, the translation engine 80 generates a fragment-context 92 for the identified translation entry. The fragment-context 92 may be any structure or data capable of associating secondary fragment 34 with address translation information for translating addressing information in the secondary fragments 34. As described in more detail below, the fragment-context 92 is used to associate secondary fragments

34 with the translation entry 90 in the translation table 82 to allow translation of the addressing information in the secondary fragments 34. In one embodiment, the fragment-context 92 is created using the 16-bit identification information in the IP header 40 of the primary fragment 32. The fragment-context 92 is associated with the translation entry 90 identified by the primary fragment 34. In this embodiment, the translation engine 80 may translate secondary fragments 34 by matching their IP header to the fragment-context 92 and using the associated entry 90 for translation.”

As is evident from the quoted passage above, once the correct translation entry 90 is identified for the secondary fragment (using the fragment-context 92 generated from the primary fragment to identify the specific translation entry 90), the translation engine 80 indeed performs an address translation for the secondary packet using the translation entry 90. That is, the translation engine 80 redundantly performs the same address translation for the secondary packet as was performed for the primary packet, in order to newly identify a destination address (such as translating an inside address to an outside address) for the secondary packet—the destination address for the secondary packet needs to be newly determined by Crow, rather than being applied from the previously determined destination address of the primary fragment. In this manner, both the primary and secondary fragments can be routed to the same destination address by Crow.

Accordingly, Crow does not meet the limitations of amended claim 1 that require “applying said determined routing information that includes said destination address to said received non-head fragment.” Explained in another way, Crow’s technique needs to perform the same address translation of both the primary fragment and the secondary fragment, using the translation entry 90 to translate one address to another address, in order to separately identify a destination address. Crow’s fragment-context 92 does not contain the destination address of the head fragment—he instead needs to newly perform address translation in order to obtain the

destination addresses for his secondary fragments—he therefore does not “apply said determined routing information that includes said destination address” as recited in claim 1.

Thus, claim 1 is allowable.

Various other amendments are made to claim 1 to provide consistent antecedent basis and/or to otherwise place claim 1 in better form.

B. Discussion of the other independent claims

Independent claim 9 as amended recites, *inter alia*, “determining routing information that includes a destination address for said head fragment” and “applying the determined routing information that includes said destination address to any corresponding non-head fragment ...” As explained previously above, Crow repeats the same processing (address translation) of his fragments in order to separately and repeatedly determine the destination address—the translation entry 90 is used by his translation engine 80 for both primary and secondary packets to perform the same (redundant) address translation to determine destination addresses.

Further, it is noted that in Crow, the fragment-context 92 is the only piece of data from the processing of his primary fragment that is used in connection with the processing of his secondary fragments. The fragment-context 92 does not and cannot include “said destination address” of the head fragment as recited in claim 9, since Crow has to newly perform a translation of the secondary fragment in order to determine its destination address.

Accordingly, claim 9 is allowable.

Amended independent claims 13 and 17, as well as new independent claim 28, contain limitations similar to claim 9 (using varying language), and are allowable for similar reasons.

Amended independent claim 20 contains language similar to claim 1 (using varying language, and is allowable for similar reasons.

C. Discussion of the dependent claims

Dependent claim 7 recites that the updating the non-head fragment with routing information includes “adding a routing tag to the non-head fragment.” Page 8 of the present Office Action admitted that Crow does not disclose this feature, but nevertheless cited Viswanath as allegedly supplying the missing teachings of Crow.

It is respectfully submitted that Viswanath does not cure the deficiencies of Crow.

Specifically, nothing is disclosed, taught, or suggested by Viswanath that his routing tag is added “to the non-head fragment.” Indeed, it appears that Viswanath instead provides his routing tag to the “switch fabric” or to his “switching logic,” and does not add the routing tag to his fragment. The fact that the routing tag of Viswanath is provided to his switch fabric/logic, rather than being added to a fragment, is explained below from his column 8, lines 36-51 (emphasis ours):

“If in step 102 the flow module 44 determines that the flow is not fragmented, then the policy filter 80 determines the policy identifier in step 106 based on the IP source address, and the TCP source port address. The policy filter determines the best match for the policy identifier in step 108, and accesses the policy cache 84 to determine whether the identified policy tag is present; if the identified policy tag is not in the policy cache 84, the policy filter 80 fetches the table entry from the policy table 28 the and stores the fetched entry into the cache 84 in step 110. The policy filter 80 then outputs the policy identifier as a template tag to the layer 3 switching logic in the switch fabric 25 in step 120.

If in step 102 the flow module 44 determines that the flow is fragmented, if the fragment is the first fragment in step 112, the policy filter 80 performs the same operation in steps 106, 108, and 110, and outputs the template tag to the switch fabric in step 120. However in this case, the flow identification module 82 monitors the output bus of the policy filter 80, and stores into the table 86 the IP source address, IP

destination address (optional), the IP identifier, and the corresponding policy tag of the first IP fragment in step 114.

If in step 112 the fragment is not the first fragment, the flow identification module 82 performs a lookup in the table 86, using the IP source address and IP identifier of the corresponding received frame fragment. The flow identification module 82 then outputs the tag for use by the layer 3 switching logic in step 120."

Accordingly therefore, it is respectfully submitted that even if Viswanath is combined with Crow, as the present Office Action has attempted to do, the limitations of claim 7 that require the routing tag to be added to the fragment would still not be met. Accordingly, claim 7 is allowable.

Dependent claims 16 and 30 also recite the routing tag added to the fragments, and thus are allowable as well.

D. Other claim amendments

Various other amendments are made to the claims as shown to provided consistent antecedent basis, to provide consistent recitations between the claims, and/or to otherwise place such claims in better form.

III. Conclusion

It is respectfully submitted that the independent claims are in condition for allowance. The dependent claims that depend directly or indirectly on these independent claims are likewise allowable based on at least the same reasons and based on the recitations contained in each dependent claim.

If the applicant's attorney (Dennis M. de Guzman) has overlooked a teaching in any of the cited references that is relevant to the allowability of the claims, the Examiner is requested to specifically point out where such teaching may be found. Further, if there are any

informalities or questions that can be addressed via telephone, the Examiner is encouraged to contact Mr. de Guzman at (206) 622-4900.

The Director is authorized to charge any additional fees due by way of this Amendment, or credit any overpayment, to our Deposit Account No. 19-1090.

All of the claims remaining in the application are believed to be allowable. Favorable consideration and a Notice of Allowance are earnestly solicited.

Respectfully submitted,
SEED Intellectual Property Law Group PLLC

/Dennis M. de Guzman/

Dennis M. de Guzman
Registration No. 41,702

DMD:wt

701 Fifth Avenue, Suite 5400
Seattle, Washington 98104
Phone: (206) 622-4900
Fax: (206) 682-6031

1000071_1.DOC